Observing to the very edge of a black hole using wideband signal processing

Jonathan Weintroub for the Event Horizon Telescope Collaboration
http://eventhorizontelescope.org/
Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138
E-mail: jweintroub@cfa.harvard.edu

Abstract. A broad international collaboration is building the Event Horizon Telescope (EHT). The aim is to test Einstein's theory of General Relativity in one of the very few places it could break down: the strong gravity regime right at the edge of a black hole. The EHT is an earth-size VLBI array operating at the shortest radio wavelengths, that has achieved unprecedented angular resolution of a few tens of μarcseconds. For nearby super massive black holes (SMBH) this size scale is comparable to the Schwarzschild Radius, and emission in the immediate neighborhood of the event horizon can be directly observed. We give an introduction to the science behind the CASPER-enabled EHT, and outline technical developments, with emphasis on the secret sauce of high speed signal processing.

Figure 1. The Submillimeter Array (SMA), an 8 antenna radio interferometer on Mauna Kea in Hawaii, is the author's home facility and a key station in the Event Horizon Telescope. The SMA operates at wavelengths of 1.3 mm and shortward. Its eight six meter dishes are on the right of this panorama of Mauna Kea's "Submillimeter Valley". The two single dish facilities on the left are the Caltech Submillimeter Observatory (CSO, the silver ball) and the James Clerk Maxwell Telescope (JCMT, the hexadecagonal cylinder).

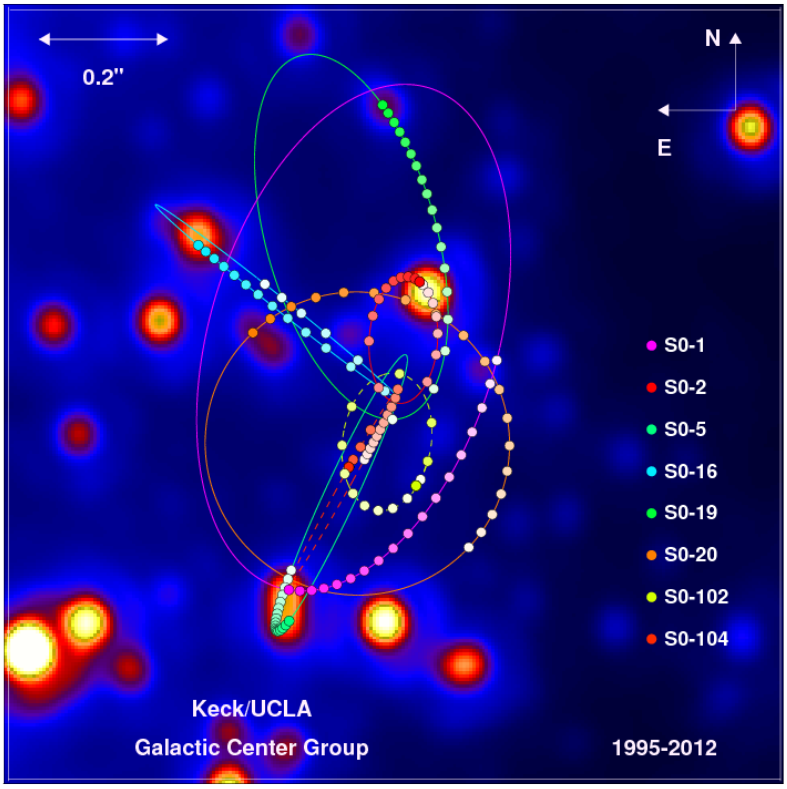

Figure 2. A diffraction-limited image is overlaid with the measured orbits of stars within the central 1.0 × 1.0 arcseconds of our Galaxy. Stars that have been observed through at least one turning point of their orbit best constrain the orbital parameters. The annual average positions for these stars are plotted as colored dots, and the best fitting simultaneous orbital solutions are lines. These orbits are the best evidence for a supermassive black hole, which has a mass of 4 million times the mass of the Sun. From [7]. This image was created by Prof. Andrea Ghez and her research team at UCLA and are from data sets obtained with the W. M. Keck Telescopes.

1. Science introduction

In the past two decades it has become widely accepted that most galaxies harbor supermassive black holes (SMBH) at their centers. The closest SMBH is associated with the radio source Sagittarius A* (Sgr A*) at the center of the Milky Way. Sgr A* was discovered by Balick and Brown in 1974 [1].

Infrared observations made with adaptive optics on the Keck 10 m telescopes on Mauna Kea (Hawaii, USA), and the European Southern Observatory (ESO) Very Large Telescope (VLT, Cerro Paranal, Antofagasta, Chile) have been used to measure the orbits of a number of stars bound to Sgr A* with periods as short as 15 years ([2], [3], [4]). From these orbits the central mass has been determined to be about $4 \times 10^6 M_\odot$. Further, the mass density must be greater than 10^{22} $M_\odot pc^{-3}$ because lack of proper motion of Sgr A* implies that $> 10\%$ of the mass must be tied to Sgr A*, which has size < 1 AU [5]. Barring a few really exotic explanations this mass is likely in the form of a black hole. The Schwarzschild radius (R_{Sch}) for this black hole is about 1.2×10^{12} cm or 10 μas at a distance of 8 kpc. The closest approach among the stellar orbits now measured is about 560 R_{Sch}. Figure 2 shows the orbits of stars in the central square arcsecond of the galaxy as measured over 17 years by the UCLA Galactic Center Group.

The spectral energy distribution of Sgr A* peaks in the submillimeter wavelength part of the spectrum. Marrone [6] using the Submillimeter Array (SMA) shows that the peak is near 0.8 mm, and further that around 1 mm wavelength the emission becomes optically thin. Hence observations made at or shorter than this wavelength offer the possibility of observing emission originating close to the event horizon. At radio wavelengths the image of Sgr A* is blurred by the turbulence of the ionized gas in the inner galaxy. The fluctuations in the index of refraction due to this turbulence broadens the image to about 15 $\times(\lambda/1\text{mm})^2$ μas. Hence the part of the spectrum near 1 mm wavelength is ideal to probe the black hole because the flux density

is high, the emission is relatively unaffected by scattering in the intervening medium, and the opacity is small enough to observe close to the event horizon. Optically, Sgr A* is obscured by dust. See figure 3 which shows the λ^2 scattering relationship together with size estimates from observations, and the transition between the "scattering dominated" and "source dominated" observational regime.

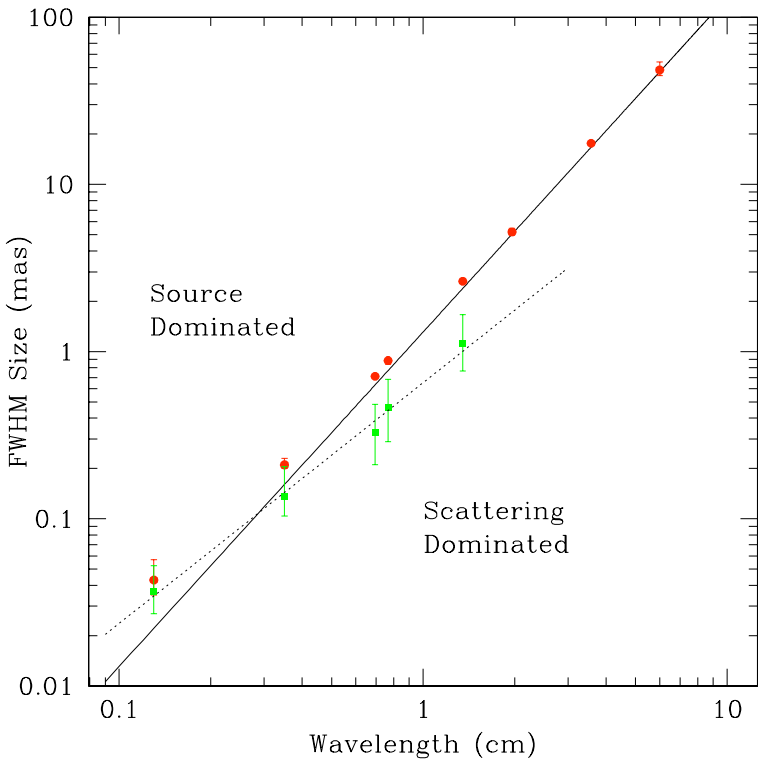

Figure 3. Observed and Intrinsic size of Sgr A* as a function of wavelength. Red circles show major axis observed sizes of Sgr A* from VLBI observations (all errors 3σ). Data from 6 cm to 7 mm wavelength are from [12], data at 3.5 mm wavelength is from [13], and the 1.3 mm data is from [10]. The solid line is the best-fit λ^2 scattering law from [12], and is derived from measurements made at $\lambda > 17$ cm. Below this line, measurements of the intrinsic size of Sgr A* are dominated by scattering effects, while measurements that fall above the line indicate intrinsic structures that are larger than the scattering size. Green points show derived major axis intrinsic sizes from 2 cm $< \lambda <$ 1.3 mm and are fit with a $\lambda^{\alpha} (\alpha = 1.44 \pm 0.07, 1\sigma)$ power law shown as a dotted line. From [10]

2. The shadow of the event horizon

When viewed from afar at spatial resolutions typically available in astronomy, a black hole appears to be exceptionally bright: a paradox created by its own intense gravitational pull. If, however, it were possible to zoom in to the event horizon and capture an image, General Relativity predicts that we would see a "shadow" against a backdrop of glowing super-hot gas as the black hole absorbs and bends the light surrounding it. Since the shadow was first predicted by James Bardeen in the 1970s, [8] direct observation of this shadow and the motion of matter as it plunges inward have been long-standing goals in astronomy and physics, because the shape of the shadow and dynamics near the black hole hold answers to some of the most fundamental questions in astrophysics:

- Do Event Horizons, one-way portals in space-time, exist?
- Was Einstein right about gravity? Do his theories hold when gravity is dominant?
- How do super massive black holes affect the evolution of galaxies?

The relativistic gas in the accreting envelope of Sgr A* is the probable origin of the emission at radio wavelengths as well as in other parts of the spectrum. Theoretical calculations by

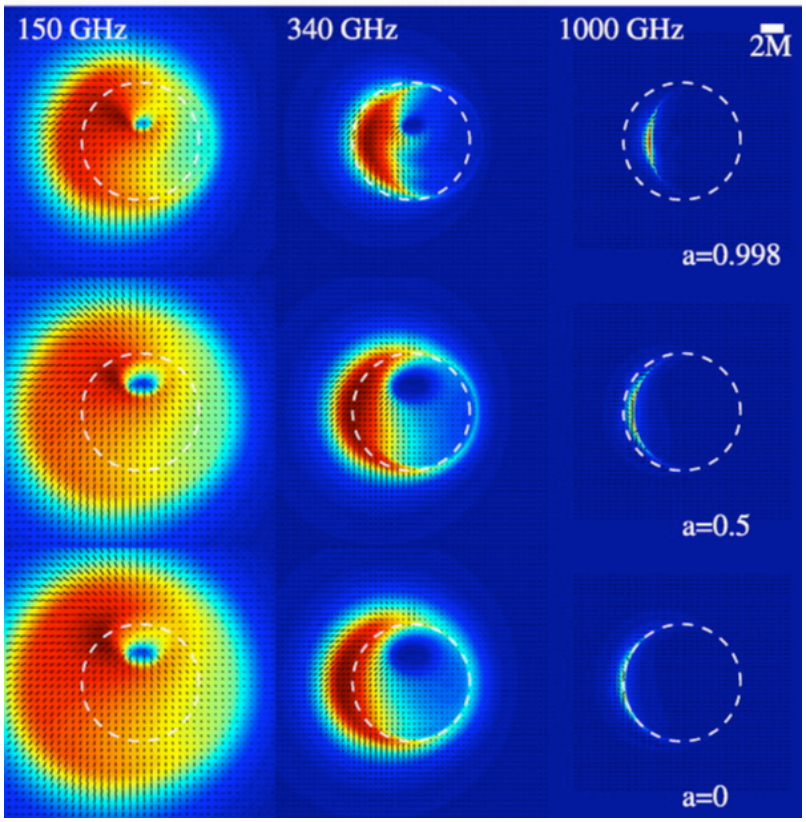

Figure 4. Theoretical calculations of images of plasma orbiting a black hole made by Broderick & Loeb ([14]). Observing frequency is indicated for each column at the top and black hole spin is indicated for each row at the right. The Schwarzschild radius (R_{Sch}) is equal to $2M$, as indicated in the upper right corner. The photon-capture radius of a non-rotating black hole is delineated by the dashed white circle. The results from our first VLBI detection at 230 GHz (2008, see [10]) are consistent with a highly time variable source of a size comparable to the dashed white circle.

Falcke, Melia and Algol ([9]) and more recently by Broderick and Loeb ([14] and see figure 4) show that the general relativistic ray paths are highly curved and lead to a severely distorted image with a central shadow. The expected image size is about 50 μas and the symmetry is sensitive to the viewing angle of the accretion disk and the spin.

For decades the possibility of making an image of the black hole shadow seemed unrealistic, because the necessary resolving power required for such work was simply not available. However, recent work has now brought this long-sought goal within reach. By refining the technique of Very Long Baseline Interferometry (VLBI), in which radio dishes separated by large geographical distances are linked together to form an Earth-sized virtual telescope, the required resolving power has been achieved. First observations carried out with three radio dishes combined in this way (in Hawaii, California, Arizona) have resulted in the exciting confirmation of event horizon size structures in Sgr A* [10] (see figure 5 and the $\sim 10^3\times$ more massive black hole in the nearby galaxy Messier 87 (M87 or Virgo A). The breakthrough refinement has been to extend VLBI to the highest observing radio frequencies ever used, where one can see clearly to the center of our Galaxy and resolving power is maximized, as well as to dramatically expand the processing and recording bandwidth of VLBI systems to achieve the needed sensitivity.

3. The Event Horizon Telescope (EHT)

The Event Horizon Telescope (EHT; `www.eventhorizontelescope.org`, [15]) is a VLBI array which uses as its stations an ad hoc combination of existing millimeter and submillimer telescopes. The EHT is designed to resolve a massive black hole with Schwarzschild Radius resolution at $\lambda \leq 1.3$ mm. With the addition of new VLBI sites, including phased ALMA in Chile, the SPT at the South Pole, and the Greenland Telescope (GLT) in Greenland, the EHT is set to grow into a true imaging array. It will reveal strong field GR effects, time resolve black

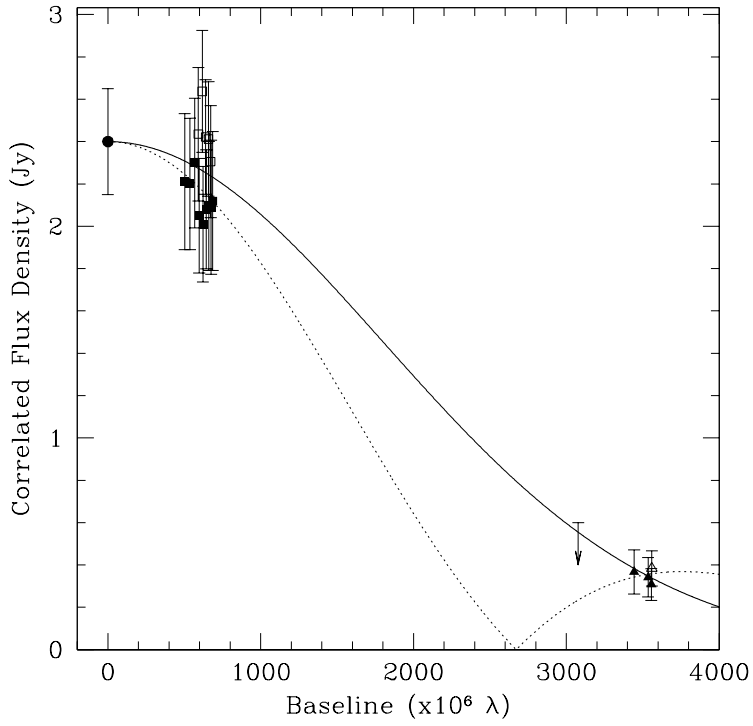

Figure 5. Fitting the size of Sgr A* using a visibility-distance plot (correlated flux density agains baseline length) and data from the 2007 campaign. Squares show Arizona-California baseline data and triangles show Arizona-Hawaii data vs. baseline length (open symbols 10 April, and filled 11 April). The best-fit circular Gaussian brightness distribution (solid line) has a FWHM size of 43.0 μas. The intrinsic size (scattering deconvolved) is estimated at 37 μas. A uniform thick-ring (i.d. of 35 μas and o.d. 80 μas) convolved with scattering is shown (dotted line). The total flux density from the CARMA array in California (2.4 Jy $\pm$ 0.25 : 1σ) is a filled circle. A 0.6 Jy flux density upper limit derived from non-detection on the Hawaii-California baselines is represented with an arrow near a baseline length of 3075×10^6. From [10]

hole orbits, and constrain models of black hole accretion and emission on the smallest scales.

M87 is giant elliptical galaxy hosting a 6 billion solar mass black hole, with Schwartzschild radius just 20% smaller in angular size compared to Sgr A*, and a prominent jet. Like Sgr A* it is a compelling target for the EHT. We have detected M87 on event horizon scales in multiple epochs [11]. For both Sgr A* and M87 the addition of new sites will lead to estimates of black hole spin and accretion disk inclination ([15], [16]).

The EHT has observed time variability of Sgr A* on a few Schwarzschild radius scales ([17]), which can be modeled as orbiting hot-spots in the surrounding accretion flow [18], [19]. The key for many of the EHT science goals is measurement of closure phase information across the array: the sum of interferometric phase around a closed triplet of EHT sites. Using closure phase techniques the EHT can detect the periodic signature of hot-spot orbits around Sgr A* with 10 second resolution [20] and provide a new method of determining black hole spin. Ultimately, the EHT will be used to make images of the black hole during periods of quiescence when strong field GR effects such as the black hole shadow can be detected. The shape and size of the shadow contains the imprint of the black hole mass and spin [21].

Three-baseline closure phase measurements have recently been demonstrated with the EHT on the calibrator quasar 1924-292 ([22] and see figure 11), indicating that with more sensitivity the array will soon be able to make similar measurements for Sgr A*. The sine-qua-non of both the closure phase variability, and ultra high angular resolution imaging observations, is improved sensitivity which is achieved by increasing the EHT bandwidth using ultra high performance signal processing.

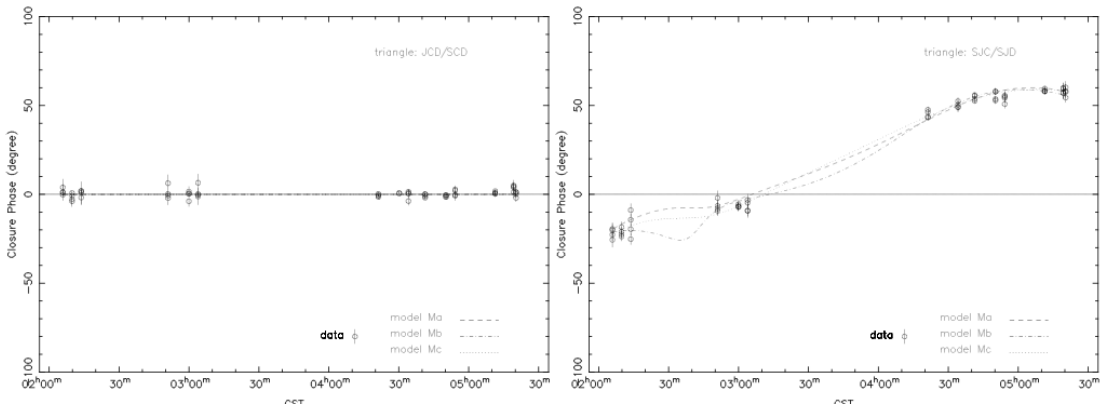

Figure 6. Plot of closure phase as a function of time for the SCD/JCD triangles (left), and SJC/SJD triangles (right) for the 3-day observations. The predicted closure phases for the models are also shown. Note that the expected closure phase for the trivial SCD/JCD triangles is zero. From [22].

The EHT is timely for several reasons.

- **Impact of the G2 Cloud:** Gillessen et al ([23]) have discovered a gas cloud of estimated mass $10^{-5}M_{\odot}$ already interacting with Sgr A*, which may last for a decade. The cloud, which is already tidally disrupted, may increase the accretion rate at the black hole by factors ranging from a few to a thousand, with potentially dramatic consequences, and extraordinary opportunities for observing the flare response of an AGN to a known stimulus. The EHT would, for example, be able to spatially resolve the emergence of a newly formed jet in Sgr A*.
- **Geographic Coverage:** The EHT relies on multiple sites to be simultaneously operational, and there is no guarantee that all critical elements of the array will be available indefinitely. Thus the window in which to carry out EHT observations of Sgr A* and M87 is finite. Figure 7 shows locations of current and anticipated EHT sites.
- **Phasing the Atacama Large Millimeter Array (ALMA):** The ALMA Phasing Project (APP) has the objective to phase up all the ALMA dishes ($\sim$ 50 dishes, each 12m diameter) for use as an exquisitely sensitive VLBI element at 1.3 mm and 0.85 mm. The advent of ALMA will be transformative, as it increases sensitivity of the EHT by an order of magnitude and improves angular resolution by a factor of two.

4. Signal processing technology for the EHT

Digital signal processing (DSP) technology from the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) [25] have been key to the success of the EHT throughout. The primary goal of CASPER is to streamline and simplify the design flow of radio astronomy instrumentation by promoting design reuse through the development of platform-independent, open-source hardware and software. CASPER technology couples the real-time streaming performance of application-specific hardware with the design simplicity of general-purpose software. By providing parameterized, platform-independent "gateware" libraries that run on reconfigurable, modular hardware building blocks, low-level implementation details are abstracted, allowing astronomers to rapidly design and deploy new instruments.

Single dish VLBI "digital back ends" (DBEs), and multi-dish VLBI phased arrays used at EHT ad hoc stations, such as the SMA, which happen to be natively array telescopes, have used

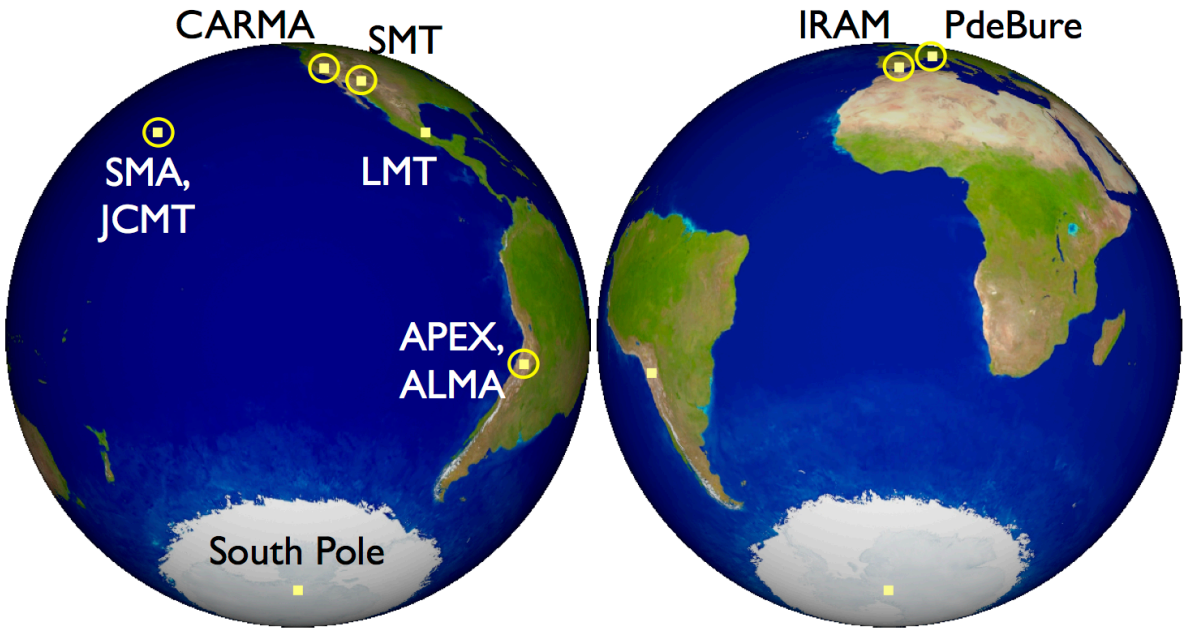

Figure 7. The geographic locations of current and future EHT sites as seen from the direction of Sgr A*. The circled locations are already outfitted for VLBI, the other locations will be commissioned within two years. The Greenland Telescope (GLT) does not have visibility of Sgr A* and is not shown.

successive generations of open source CASPER FPGA processing platforms. These include the "Internet Breakout Board" (iBOB), the second generation "Berkeley Emulation Engine" (BEE2), as well as the "Reconfigurable Open Architecture Computing Hardware first and second generations (ROACH1 and ROACH2). High speed digital signal processing is also critically dependent on high speed analog-to-digital converters, of which CASPER has a variety of open source hardware designs.

4.1. Wideband analog-to-digital converters

In a sampled data system the width of a single block of processed bandwidth is set by the ADC sample rate. If the block is narrow, the processor must be preceded by an IF system with many channels; a so called hybrid implementation, which, if large, is prohibitive in cost. Note that for an array the number of physical IF channels, the full bandwidth divided by the sampled block width, is multiplied by the number antennas $\times$ two polarizations per antenna. It is also important to appreciate that while digital technology becomes exponentially more economical with time per unit processing power, according to Moore's law, Moore's law does not apply to analog IF electronics, which typically increases in cost with time. Thus it more forward thinking and economical to design with using ADCs which are as fast as possible. Please see table 1 for a listing of various interesting ultra-wideband ADCs, including selected specifications.

Many of these achieve high sample rates by interleaving multiple slower ADC cores. Distortion results from misalignment in offset, gain and phase of the cores, as well as non-linearity . The 5 GSa/s e2v EV8AQ160 has been studied in depth [27] and has been selected for the development of a new 4 GHz correlator with integrated VLBI phased array, see section 4.2. This is a quad core device, using four 1.25 GSa/s ADC cores interleaved to achieve the 5G Sa/s design rate. The device provides register controls to align the cores to reduce the impact of spurs which arise

Table 1. Table of ultra high speed analog-to-digital converter chips, the last three entries are inexpensive comparators, essentially 1-bit ADCs.

f_s (GSa/s)	BW (GHz)	# bits	Manuf.	Part #	∼cost
5	2.0	8	e2v	EV8AQ160	$300
10	20	4	Adsantec	ASNT7120-KMA	$800
12.5	8	8	Maxtek	—	$17k
20	8	5	e2v	EV5AS210	$7k
20	13	8	Agilent	—	
20	10	3+ oflow	Hittite	HMC5401LC5	$5k
30	14	6	Micram	ADC30	$10k
56	15	8	Fujitsu	CHAIS	—
20	10	1	Hittite	HMC874LC3C	$40
12.5	14	1	Inphi	1385DX	—
25	18	1	Inphi	25707CP	—

due to misalignment in offset, gain, phase (OGP), or threshold Integral Non-Linearity (INL). All the cores are clocked by the same external clock input. In the single-channel mode, the in-phase 1.25 GHz clock is sent to ADC core A and an inverted 1.25 GHz clock is sent to core B. The 90 and 270 degree phase shifted clock signals are sent to cores C and D respectively. This interleaved mode has an equivalent sampling frequency of 5 GSa/s. A CASPER compatible printed circuit board is available based on this ADC, designed by Jiang et al. [28]. Reference [27] details how the OGP and INL corrections are derived, and figure 8 shows the improvement in Spurious Free Dynamic Range (SFDR) and Signal to Noise and Distortion (SINAD) which can be obtained using these corrections.

Looking ahead, we are presently developing a PCB based on the 26 GSa/s Hittite HMCAD5831 3-bit single core ADC with an FMC (Vita 57.1) standard interface to the Xilinx VC709 evaluation board. The VC709 is based on the XC7VX690T FPGA which includes 80 GTH multi-gigabit serial transceivers, ten of which are brought out to the FMC. The GTH transceiver is a sophisticated subsystem capable of transmitting and receiving serial data at rates up to 13.1 Gb/s per channel. Both transmitter and receiver have functional blocks essential for reliable transmission of data across this interface—these necessary features are not available in prior generations of SERDES transceiver, for example those on the Xilinx Virtex 6 family which is used on the CASPER ROACH2.

4.2. SMA Wideband Astronomical ROACH2 Machine

Table 2. Table of SWARM specifications

Feature	Specification	Remarks
Number of antennas	8	2 receivers each.
Bandwidth per receiver	2 GHz	Dual polarization in each side band.
Number of sidebands	2	90-270 Walsh splits SBs, Rx are DSB
Simultaneous receivers	2	Dual frequency or dual polarization 230 & 345 GHz
Baselines	56	28 per Rx, full Stokes, 112 total
Spectral resolution	140 kHz	2.3 GHz Nyquist / 16384 channels
Fastest dump rate	0.65 s	Single full Walsh cycle
Phased array bandwidth	4 GHz	2 GHz ×dual pol.

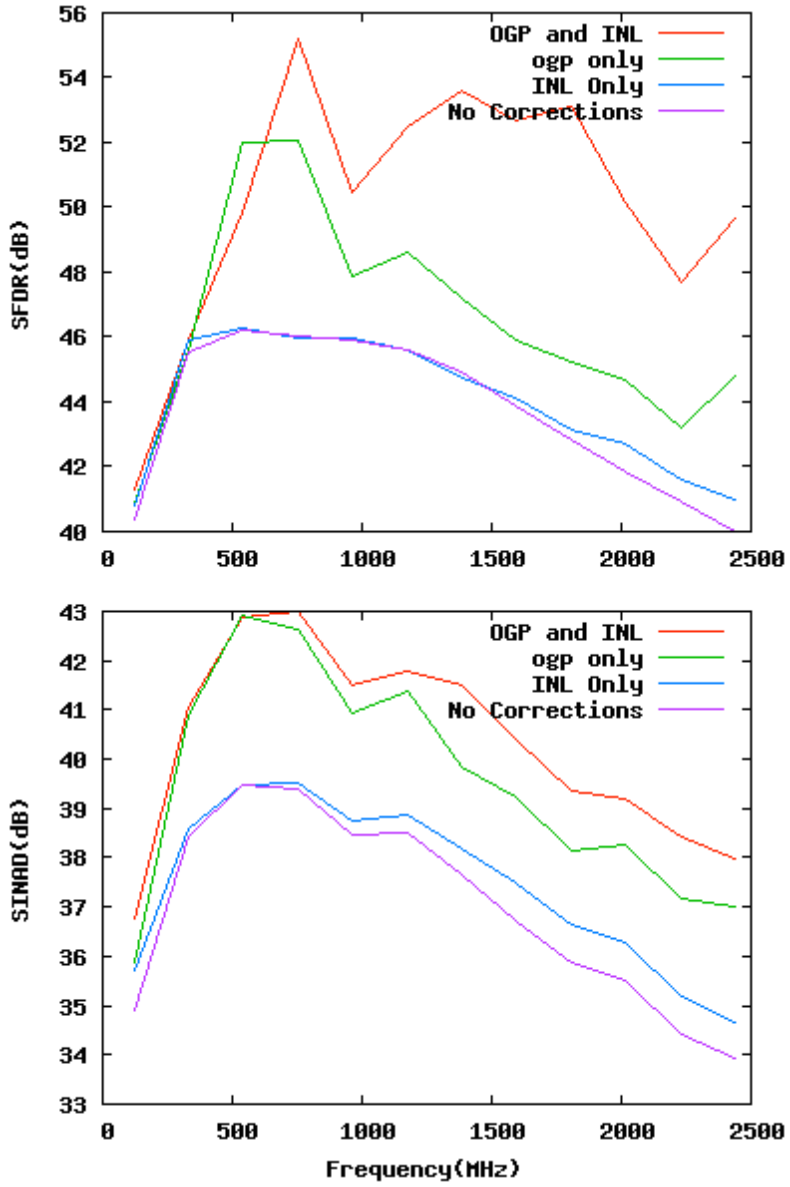

Figure 8. Spurious Free Dynamic Range (SFDR) and Signal to Noise and Distortion (SINAD) measurements as a function of frequency for the EV8AQ160. The various color coded curves show the improvement due to core calibration. Offset, gain, phase is abbreviated OGP, and integral non-linearity corrections are applied independently, and then together. It is interesting to note that INL makes only a marginal improvement when OGP is not applied, but a fairly substantial improvement when relative to a device with OGP correction already applied.

A new wide-band high spectral resolution Digital Back End (DBE) is being designed for the Submillimeter Array (SMA). The new DBE is designated SWARM for SMA Wideband Astronomical ROACH2 Machine. Figure 9 is a block diagram of the system, and table 2 gives a summary of the specifications. SWARM combines the functions of a connected element correlator, and has a natively integrated phased array data formatter for VLBI. Connected element requirements are for uniform spectral resolution of 140 kHz over a 2 GHz usable dual polarization band, to support fast spectral line surveys. The 8-element array has 28 baselines per receiver. With full Stokes capability, and including autocorrelations, the DBE must process correlations on an effective 128 baselines. The enormous computational requirements demand the latest technology, including highly parallel signal processing engines and very wide-band high sample rate analog-to-digital converters (ADCs). The requirements of SWARM are well matched to the requirements of the 5 GSa/s EV8AQ160 ADC from e2v, and this ADC was selected as the digitizer for SWARM. The CASPER ROACH2 is the processing platform.

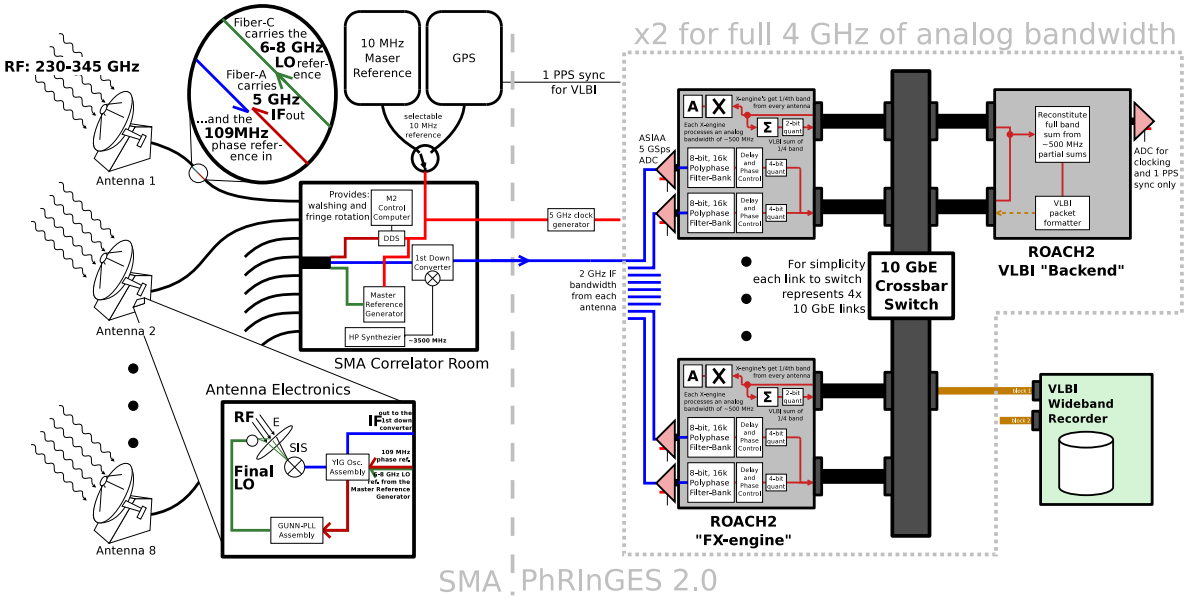

Figure 9. Block diagram of the SMA Wideband Astronomical ROACH2 Machine (SWARM)

4.2.1. FPGA DSP utilization Before committing to the ROACH2 board we needed to understand that our design would fit the Virtex 6 FPGA. The utilization is dominated by the polyphase filter bank (PFB) based on the fast Fourier transform (FFT). Consider a PFB with N points, and note that the FPGA cannot run nearly as fast as an ultra high speed ADC. A number of parallel streams must thus be processed in the FPGA, and the ratio of ADC sample rate to FPGA clock rate is the *demultiplex factor*, D.

The CASPER FFT implementation assumes that $N > D$ and that both N and D are powers of 2. It's important to note that the pipelined stages are presented only with D samples at once meaning that it takes N/D clock cycles to perform a complete computation. The subsequent FFT however, being of size D, performs a full computation per clock. This means that the full core is "streaming"; for each input time-domain sample there is exactly one output frequency-domain sample. The number of instantiated multipliers is given by equation 1 and the number of adders is given by equation 2 which have similar form. The three terms are annotated, and for most applications the FFT term dominates. Equations 1 and 2 are derived by Primiani et al. [26].

$$M_{\mathrm{PFB}} = \underbrace{D \log_2 ND}_{\mathrm{FFT}} + \underbrace{TD}_{\mathrm{FIR}} - \underbrace{2D}_{\mathrm{optimization}} \tag{1}$$

$$A_{\mathrm{PFB}} = \underbrace{\frac{3}{2} D \log_2 ND}_{\mathrm{FFT}} + \underbrace{D(T-1)}_{FIR} + \underbrace{D}_{\mathrm{reorder}} \tag{2}$$

Within a polyphase filter-bank the multiplier and adder utilization grow significantly with the demux factor, namely as $D \log_2 ND$, whereas the amount of required memory depends critically, and linearly, on the total channels, N. Generally this implies that designs with modest bandwidth but requiring significant spectral resolution will be constrained by memory. SWARM, however, has both very large bandwidth and substantial PFB size to achieve fine spectral resolution. Thus it has been necessary to find the appropriate combination of parameters

which meet the requirements of bandwidth and spectral resolution while fitting the logic and memory available in a particular FPGA. This constrains the width of the processed bandwidth block for a given FPGA, and thus provide an independent constraint on the ADC rate.

4.2.2. Early SWARM results Early tests with the SWARM correlator were run in 1/2 bandwidth mode. We tuned to the H30 α maser line in MWC349 and obtained strong fringes on all 21 baselines. Figure 10 shows the amplitude and phase spectrum for the recombination maser line. Figure 11 show the closure phase on 3C454.3 for the antenna 1-2-4 triangle of baselines. Because the source is unresolved we expect the closure phase to be zero on all triangles. The ROACH2 FPGA now has a complex bit code fully capable of connected interferometry, the remaining design challenge is to double the clock speed of the FPGA while maintaining timing closure. This step is needed for SWARM to support its full targeted bandwidth.

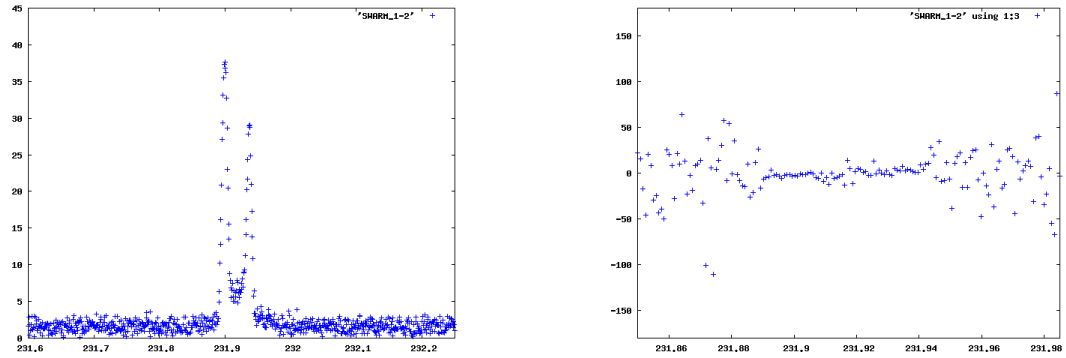

Figure 10. SWARM commissioning sky data, spectrum of the H30α maser line in the source MWC349. Left panel is cross-power amplitude, right panel is phase.

4.3. Fast digital recorders
The data rates and volumes required for EHT work are staggering. Each radio dish in the EHT can provide a radio spectrum bandwidth of 8 GHz, which is equivalent to a data rate of 8 megabytes per second. For all nine EHT sites, this corresponds to a total data rate of 36 billion bytes per second, yielding approximately 800 trillion bytes of data over a single night of EHT observations. To record and process this torrent of information requires cutting-edge systems that far exceed any VLBI instrumentation currently available. The data streams from each EHT site are recorded in real-time on Mark6 VLBI data recorders utilizing banks of hard disk and solid-state disk drives operating in parallel. These recorders must be robust against the failure or slow-down of a single drive, and this requires development of adaptive algorithms for hard disk management. The new recorder systems will achieve the same 16Gb/s data capture rates as the backends and phased array processors. Figure 12 shows photos of current Mark6 main frame and disk pack hardware.

5. Operations
Presently to mount an EHT campaign we propose independently to the participating telescopes for a scheduled window of observing time. Teams of VLBI experts travel to all EHT sites for several weeks to complete experimental setup and checkout of the VLBI equipment. Experience has shown that the dedicated observing windows do not guarantee suitable 1.3 mm weather at all EHT sites, which is essential to collect high quality data. This problem will only become

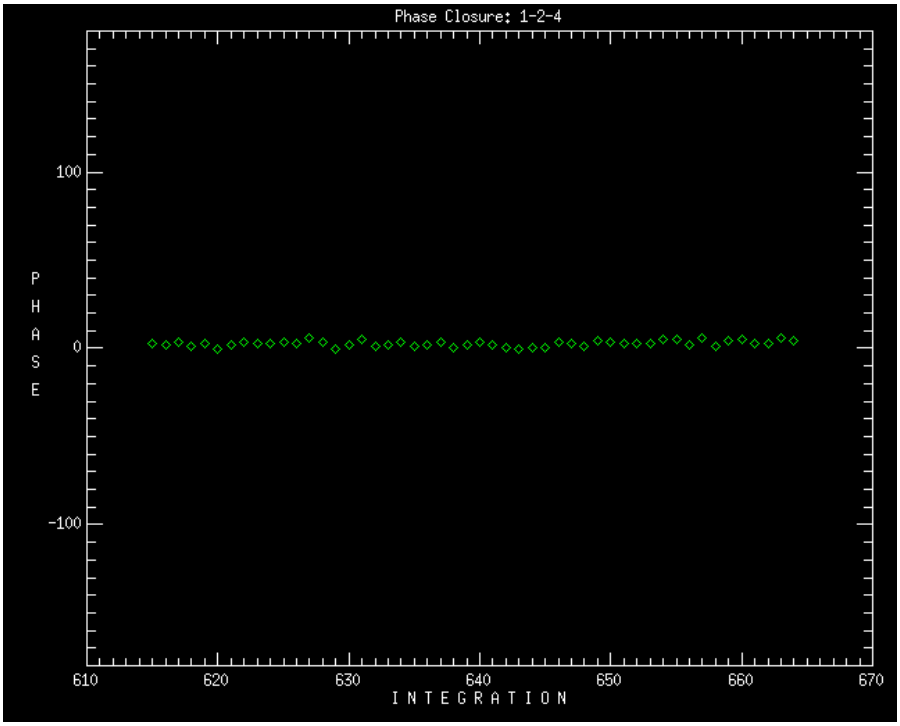

Figure 11. Closure phase on antenna 1, 2, 4 triangle

worse as the EHT array grows. A critical technical goal of EHT development is to establish the ability at each site to quickly and automatically switch to VLBI operations and carry out EHT observations remotely. When considering operations at the remote Greenland Telescope (GLT) and the South Pole Telescope (SPT) the coherence monitoring system has increased relevance, given that it is impractical to travel to these sites for specific campaigns.

5.1. Weather

Submillimeter astronomy, unlike other radio wavebands, demands excellent weather, typically obtained only at high, dry localities. The key measures are total precipitable water vapor in a column of atmosphere towards the zenith (abbreviated PWV, measured with a water vapor radiometer, and as low as $\sim$1 mm for excellent submillimeter opacity), and atmospheric stability, akin to seeing. Maximizing mutual visibility time means that for many sites observations are at low elevations through many equivalent zenith layers of atmosphere.

In VLBI the probability of good weather at all sites is multiplicatively less likely than for a single site. Of the sites which have been routinely used for EHT operations, the SMA site on Mauna Kea has the most reliable weather. In May 2012 we obtained a first fringe to the Atacama Pathfinder Experiment (APEX) telescope in Chile, also an excellent submillimeter site. It should be noted that even on Mauna Kea the stability deteriorates after sunrise, and observations are typically run at night, given time differences between the sites, the problem of good weather at the particular time of day of observations is still greater. Mauna Kea is rare in having a dedicated weather service, the Mauna Kea Weather Center. Weather predictions tuned for astronomy and specific to the actual observatory locations are considerably harder to obtain at other sites. The only scalable model is to support a turnkey "target of opportunity" mode of operations. Whenever there is a coincidence of good weather at all EHT sites, an observation must start at the flick of a switch.

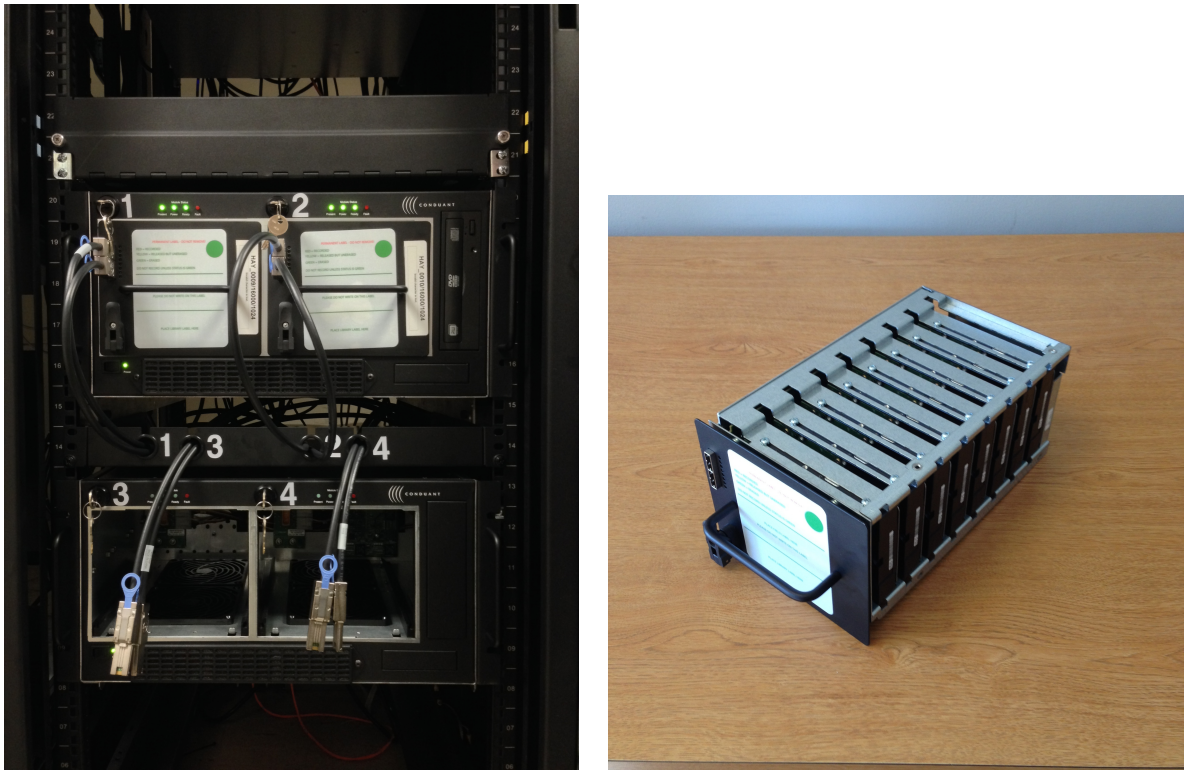

Figure 12. The left panel shows a photo of a Mark6 data recorder, and the right panel a photo of a disk pack, or "8-pack" which might currently carry eight 6-terabyte commercial hard disk drives for 48 TB of storage.

5.2. Features for turnkey EHT operation

There are two key systems designed to remotely and continuously monitor independent measures of station coherence. The first is the absolute quality of the station reference signal provided by the hydrogen maser time standard, specifically its 10 MHz phase noise, and its stability (Allan Variance). The second is a front end pure tone injection system which measures the ability of the heterodyne receiver to receive and return to VLBI data storage a phase coherent IF.

5.3. Measurement of maser stability and phase noise

In figure 13, a TSC-5115A Allan Variance test set referenced to an ultra high quality Oscilloquartz crystal is used to measure the maser's Allan Variance. Such a crystal is as stable as the maser over time periods up to 1 second. The Oscilloquartz 8607-BM 5 MHz with both Option 8 and Option L is the best device on the market. Option L is key, providing phase noise better than -130 dBc at 1 Hz. Note that maser phase noise is magnified in its effect linearly with to the ratio of the ∼230 GHz local oscillator to the maser reference frequency. This is typically over five orders of magnitude. The TSC-5115A also makes measurements of phase noise of the maser input.

5.4. Tone injection test of local oscillator and IF systems

Even with the quality of the maser reference as a given, there is still a danger that the telescope local oscillator chain, which effectively multiplies the 5 or 10 MHz from the maser more than five orders of magnitude, may corrupt the receiver reference.

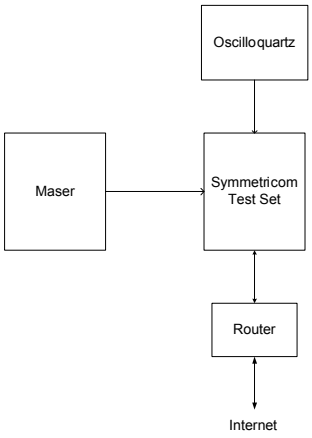

Figure 13. Setup for remote measurement of maser phase noise and Allan Variance referenced to a high quality quartz crystal from Oscilloquartz. The absolute stability of a maser used for submillimeter VLBI for 10 second integrations should be less than about two parts in 10^{14}. Symmetricom, Inc, manufacture a line of low noise phase noise measurement systems. The model TSC-5115A, which has an input frequency range from 1 to 30 MHz, is suitable for maser reference measurements whose output oscillators typically run at 5 or 10 MHz.

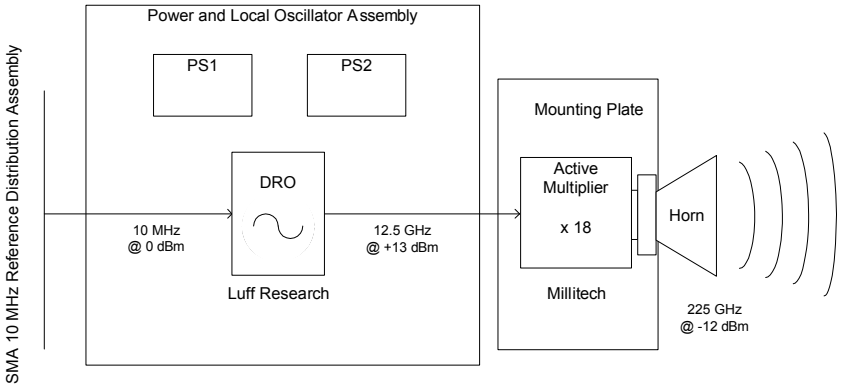

Figure 14. System for injection of a high quality tone at RF frequency at EHT stations. A Dielectric Resonant Oscillator (DRO) generates a pure tone at 12.5 GHz. This tone is multiplied in a Millitech frequency multiplier chain by a factor of 18.

In figure 14, a test tone locked to the maser reference at near to the LO frequency is coupled via an RF horn directly into the telescope's submillimeter receiver. After passing through the receiving intermediate frequency (IF) system, including fiber optic transmission, and being mixed to a much lower frequency, the signal can be directly compared to the maser in the control building. This test is independent of the test of absolute maser stability and is a comprehensive test of LO chain and IF system.

6. Summary

The EHT has showns that 1.3 mm VLBI with a 3-baseline array can probe Event-Horizon-Scale structures in Sgr A* and M87. Improvements in phased arrays, bandwidth and number of stations are expected to lead to an image of a black hole "shadow" within about two years. If detected the shadow will be compelling evidence for the existence of an event horizon. The ALMA Phasing Project (APP) will be the biggest single advance in capabilities. The EHT is a timely challenge given the state of technology, submilimeter facilities, and the G2 cloud impact. High speed signal processing, described in part in this paper, is absolutely crucial to the Event Horizon Telescope achieving its lofty goals.

7. Acknowledgments

Travel support from the HPSPSA conference organizers is very much appreciated.. The EHT receives funding from the US National Science Foundation (NSF) and the Gordon and Betty Moore Foundation (grant GBMF3561). Xilinx has donated Virtex 6 FPGAs and design software. The contributions of the open source CASPER community are gratefully acknowledged. The Submillimeter Array is a joint project between the Smithsonian Astrophysical Observatory and the Academia Sinica Institute of Astronomy and Astrophysics.

8. References

[1] Balick,B. & Brown,R.L. 1974, ApJ, 194, 265
[2] Schödel, R, et al. 2002 Nature, 419, 694
[3] Genzel, R. et al. 2003, Nature, 425, 934
[4] Ghez, A. M. et al. 2005, ApJ, 620, 744
[5] Reid, M. J. & Brunthaler, A. 2004, ApJ, 616, 872
[6] Marrone, D., Moran, J. M., Zhao, J.-H., and Rao, R. 2007, ApJ, 654, L57
[7] `http://www.galacticcenter.astro.ucla.edu/pictures/orbitsOverImage12.shtml`
[8] Bardeen, J. M. 1973, Black Holes, ed. C. DeWitt & B. S. DeWitt (New York: Gordon and Breach), 215
[9] Falcke, H., Melia, F.,& Algol, E. 2000, ApJ, 528, L13
[10] Doeleman, S. S., Weintroub, J., Rogers, A. E. E., Plambeck, R., Freund, R., Tilanus, R. P. J., Friberg, P., Ziurys, L. M., Moran, J. M., Corey, B., Young, K. H., et al. 2008, Nature, 455, 78
[11] Doeleman, S. S., et al., 2012 Science, **338**, 355D
[12] Bower, G. C., Falcke, H., Herrnstein, R.M., Zhao, J.H., Goss, W.M. & Backer, D.C. 2004, Science, 304, 704
[13] Shen, Z. Q., Lo, K. Y., Liang, M. -C., Ho, P. T. P. and Zhao, J.-H., 2005, Nature, 438, 62
[14] Broderick, A. E. & Loeb, A. 2006, MNRAS, 367, 905
[15] Weintroub, J. 2008 "A submillimeter VLBI array", *Proceedings of The Universe Under the Microscope: Astrophysics at High Angular Resolution*, Bad Honnef, Germany, April 2008, Editors: R. Shödel, A. Eckart, S. Pfalzner & E. Ros, Journal of Physics Conference Series
[16] Broderick, A. E., Fish, V. L., Doeleman, S. S. & Loeb, A. 2011b, ApJ, 738, 38
[17] Fish, V. L., Broderick, A. E., Doeleman, S. S. & Loeb, A. 2009, ApJL, 692, 14
[18] Broderick, A. E., Fish, V. L., Doeleman, S. S. & Loeb, A. 2009, ApJ, 697, 45
[19] Broderick, A. E., Fish, V. L., Doeleman, S. S. & Loeb, A. 2011a, ApJ, 735, 110
[20] Doeleman, S. S., Fish, V. L., Broderick, A. E., Loeb, A., & Rogers, A. E. E. 2009, ApJ, 695, 59D
[21] Johannsen, T. & Psaltis, D. 2010, ApJ, 716, 187
[22] Lu, R. et al, 2012, ApJL, 757, 14
[23] S. Gillessen, R. Genzel, T. K. Fritz, E. Quataert, C. Alig, A. Burkert, J. Cuadra, F. Eisenhauer, O. Pfuhl, K. Dodds-Eden, C. F. Gammie & T. Ott', Nature **481**, 5154
[24] Fish, V. L., et al., 2011, ApJ, **727L**, 36f
[25] `http://casper.berkeley.edu`
[26] `https://www.cfa.harvard.edu/twpub/SMAwideband/MemoSeries/sma_wideband_utilization_1.pdf`
[27] N. A. Patel, R. W. Wilson, R. Primiani, J. Weintroub, J. Test, K. Young, 2014, Journal of Astronomical Instrumentation, 3, 1
[28] Jiang, H., Liu, H., Guzzino, K. et al., 2012, Digitizing the Yuan Tseh Lee Array for Microwave Background Anisotropy by 5 Gsps ADC boards, In *IEEE International Conference on Electronics, Circuits and Systems*, p. 304.